Operational Evidence and Outcomes
Volume 1 · Number 1 · 2026

Inverting a Pose Does Not Invert a Camera Path: A Coordinate-Convention Reproduction Audit

Sean Robertson[1] · Ethan Butler[2] · James Harrison[3]

[1] Belgium · [2] Austria · [3] Italy

Original publication: November 8, 2026 | Article 689

Abstract

A translation column is not always a camera position. Reading the translation of inverse KITTI poses as world positions produced 95.08 km of apparent travel, compared with 22.18 km from the correct camera centers. Adding a common 100-m offset to world coordinates left the correct distances unchanged but raised the mistaken total to 103.38 km. Across eleven sequences, mean relative overstatement was 255.74%. The experiment illustrates an implementation-convention failure, without attributing that failure to a published method or treating it as physical camera motion.

Deliberately incorrect pose extraction compared with a correct geometric interpretation. The public dataset is not alleged to contain the error demonstrated here.

Inverting a Pose Does Not Invert a Camera Path: A Coordinate-Convention Reproduction Audit

> Deliberately incorrect pose extraction compared with a correct geometric interpretation. The public dataset is not alleged to contain the error demonstrated here.

Abstract

A translation column is not always a camera position. Reading the translation of inverse KITTI poses as world positions produced 95.08 km of apparent travel, compared with 22.18 km from the correct camera centers. Adding a common 100-m offset to world coordinates left the correct distances unchanged but raised the mistaken total to 103.38 km. Across eleven sequences, mean relative overstatement was 255.74%. The experiment illustrates an implementation-convention failure, without attributing that failure to a published method or treating it as physical camera motion.

1. A reproducible file can still be interpreted incorrectly

A camera-pose file can be parsed without an exception, inverted correctly as a matrix, and still be used incorrectly in a movement calculation. The difficulty is that a transformation's translation component depends on the direction in which coordinates are mapped. A camera-to-world matrix and its world-to-camera inverse describe the same geometry, but their last columns are not interchangeable camera positions. A path assembled from the wrong column can look numerically well formed while measuring a different object.

We ask how much that deliberate convention mistake changes path-distance reporting in actual public reference trajectories. We also ask whether the mistaken result depends on a physically irrelevant choice of world-coordinate origin. The second question provides an especially clear diagnostic: changing a common coordinate origin should not change the length of the correctly represented camera-center path.

Meng, 2026 considers scale-calibrated camera-motion transfer through a gauge-invariant parallax quantity. That work motivates specifying coordinate conventions before interpreting a geometric scalar. The current study does not compute its parallax number, reproduce its transfer procedure, or evaluate generated motion. It examines a prerequisite for such measurements: extracting the intended trajectory from a transformation array consistently.

The tested propositions are concrete. A correct inverse representation should recover the original camera centers after the appropriate conversion. A deliberately naive inverse-column extractor may yield a different path length. Furthermore, a common translation of world coordinates may change that erroneous length because the inverse translation is expressed through the camera orientation. These are computational consequences checked against reference arrays, rather than claims about an estimator's real-world error rate.

2. Geometry methods share matrices but require explicit directions

Multi-view geometry Hartley and Zisserman, 2004 distinguishes coordinate transformations from the physical entities described by them. Pose-estimation procedures such as EPnP Lepetit et al., 2009 and mapping systems such as ORB-SLAM2 Mur-Artal and Tardos, 2016 operate within specified frame conventions. Structure-from-motion pipelines such as COLMAP Schonberger and Frahm, 2016, bundle-adjustment analysis Triggs et al., 2000, and sparse bundle adjustment Lourakis and Argyros, 2009 provide further contexts where a consistent mapping direction is essential. Their use of transformation matrices does not make every matrix translation a world camera center.

Driving datasets provide many opportunities to encounter such distinctions. KITTI Geiger et al., 2012 supplies the reference poses used here. Virtual KITTI Gaidon et al., 2016 and Virtual KITTI 2 Yohann et al., 2020 offer other scene and camera settings. nuScenes Caesar et al., 2019, the Waymo Open Dataset Sun et al., 2019, and Argoverse 2 Wilson et al., 2023 organize additional driving observations. We do not pool these sources or compare their coordinate conventions empirically. Their relevance is the need to carry representation metadata when geometry moves between data-processing systems.

Rotation is itself a structured object. Quaternion-based animation Shoemake, 1985 and circular-data analysis Fisher, 1993 motivate preserving orientation geometry rather than treating angles or matrix entries as arbitrary Euclidean channels. In the present experiment, however, no orientation interpolation or circular regression is performed. Orientations enter only through proper rigid-transform inversion and the deliberate misuse of the resulting translation column.

Measurement-scale analysis Stevens, 1946 also helps identify the claim at stake. A distance in meters is interpretable only if its input coordinates refer to positions in a common frame. Printing the unit next to a derived number does not validate that requirement. Measurement-error reasoning Carroll et al., 2006 and agreement analysis Bland and Altman, 1986 consequently motivate checking the geometric object before comparing the magnitude of two numerical summaries.

3. Source trajectories and the correct camera center

The local KITTI pose archive contains eleven files and 23,201 three-by-four pose rows. We include all available reference sequences, with no selection based on the subsequent error size. The official benchmark page links the pykitti tools, whose odometry loader identifies these pose matrices as camera-to-world transformations Clement, 2026. We retain the downloaded source file, its checksum, and the loader snapshot used to inspect that convention.

Let the leading rotation block map camera coordinates into the world frame, and let the last column be the camera center in that world frame. The inverse transformation has the transposed rotation and a translation equal to the negative transposed rotation applied to the original camera center. That inverse translation locates the world origin in camera coordinates. It is not the original camera position expressed in the world frame.

The correct center can be recovered from the inverse by applying the negative original rotation to its translation. We implement this round trip for every source pose. Before inversion, rounded source rotation matrices are mapped to proper rotations using SciPy's rotation construction Virtanen et al., 2020. This makes the forward and inverse operations geometrically consistent rather than relying on slight text-rounding departures from exact orthogonality.

The maximum absolute coordinate difference in the proper round trip is approximately 1.14×10^{-12} meters. That number describes numerical agreement in the implementation, not physical accuracy of the

source trajectory. Numerical-analysis guidance Higham, 2002 and floating-point principles Goldberg, 1991 motivate keeping that distinction explicit. The rounding residual is tiny compared with the path discrepancies generated by the deliberate convention error.

4. The counterfactual error and an origin check

The incorrect extractor takes each inverse matrix's translation column and treats it as a world camera position. It then sums Euclidean distances between consecutive extracted vectors. This operation mixes camera-dependent coordinate frames across observations. We introduce it intentionally as a reproducibility failure mode; it is not an allegation about the KITTI archive, its loader, or any cited estimation method.

The origin check changes the world coordinates so that every correct camera center receives the same offset of one hundred meters on the first coordinate axis. Rotation blocks remain unchanged. We then invert these consistently transformed poses and rerun both center extraction rules. The physical relative movement is unchanged. A correct path-distance calculation should consequently be invariant to this common translation of coordinates.

The largest sequence-level change in correct path length is approximately 2.27×10^{-13} meters. The proper inverse recovery therefore passes the origin diagnostic to numerical precision. The naive inverse-column path has no corresponding invariance guarantee because its added term depends on orientation at each pose. A common coordinate offset can become a varying vector in the camera-dependent representation.

Figure 1 separates the correct geometric round trip from the deliberately wrong extraction. It also marks the origin transformation as a change in representation, rather than a second physical journey. This distinction prevents a larger resulting path sum from being described as additional travel by the camera or as an observed environmental change.

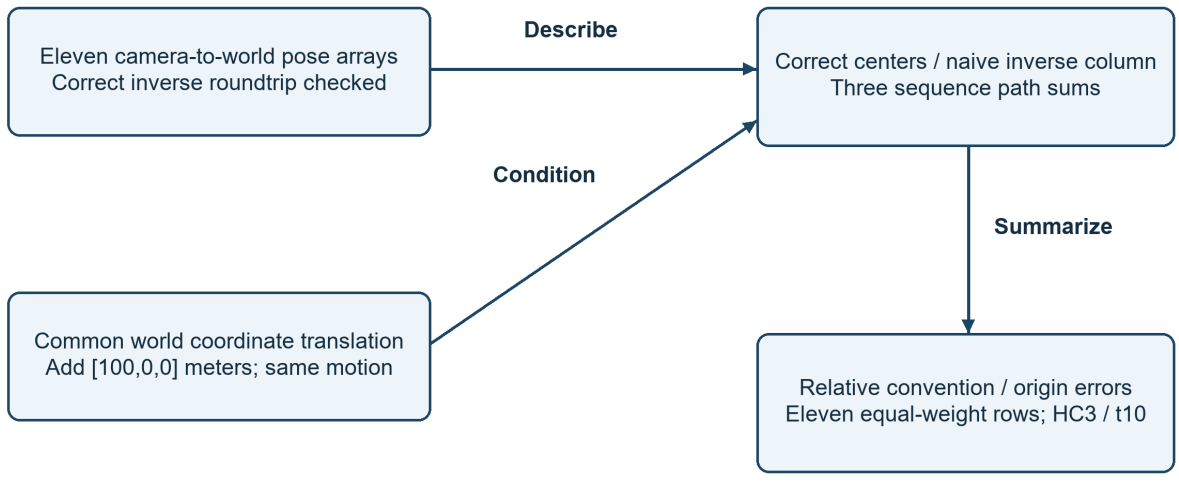

Figure 1. *Source camera centers and correctly recovered inverse centers describe the same path. The inverse translation column alone is a different vector, whose apparent path can change with the world origin.*

5. Sequence-level outcomes and regression descriptions

For each sequence, the reference distance is the polyline length through consecutive source camera centers. The first outcome is naive inverse-column distance divided by reference distance, minus one. It is a relative reporting error: one means a one-hundred-percent overstatement. The second outcome is the origin-shifted naive distance minus its original naive distance, divided by the same reference distance. A positive value indicates additional error introduced solely by the coordinate change.

We summarize each outcome with an intercept-only OLS model giving equal weight to the eleven sequences. HC3 standard errors White, 1980 and t intervals use ten residual degrees of freedom. Clustered-inference guidance Cameron and Miller, 2015 cautions that the reference recordings do not form a random sample of independent road environments. The intervals are working descriptions of heterogeneity in a selected benchmark cohort, not uncertainty in a universal probability of making this programming mistake.

The observations are sequences rather than the 23,201 poses, because each sequence supplies a full path-distance ratio. Treating pose rows as independent replications of a trajectory-level error would inflate the nominal sample size. The source's adjacent camera centers remain a captured reference, not a continuous physical-motion truth. Motion between source observations and reference estimation error are outside the experiment.

Concordance reasoning Lin, 1989 is relevant to a separate interpretation boundary. Recovering the correct center numerically is an agreement check between equivalent representations. The regression contrasts then describe a knowingly nonequivalent extraction. They do not estimate how well a new pose estimator agrees with the reference, and they are not official KITTI odometry translation or rotation error scores.

6. Results: a large discrepancy from a small metadata omission

The total reference path length across the eleven sequences is 22,179.67 meters. The naive inverse-column calculation gives 95,077.85 meters, even though it begins with inverses of the same source poses. After the common coordinate offset, the mistaken total rises to 103,383.50 meters. These are archive totals across separate recordings, not three journeys observed in the world.

Table 1. Relative path-reporting errors from inverse-column misinterpretation.

Model	Term	Estimate (SE)	95% CI	N
M1	Mean outcome/contrast	2.557380 (0.464061)	[1.523387, 3.591373]	11
M1: fit	R-squared = 0.0000	-	-	11
M2	Mean outcome/contrast	0.335356 (0.127122)	[0.052111, 0.618602]	11
M2: fit	R-squared = 0.0000	-	-	11

Model labels: M1 = Inverse-column path error relative to camera-center path; M2 = Additional relative error after a common 100-m world-origin shift.

Model 1 estimates mean relative overstatement at 2.557380, or 255.74 percent, with standard error 0.464061. Its 95% interval is [1.523387, 3.591373], corresponding to 152.34-359.14 percent. Model 2 estimates an additional origin-induced relative error of 0.335356, with interval [0.052111, 0.618602]. Both models have eleven observations, no controls, HC3 uncertainty, and effectively zero R-squared.

The mean relative overstatement should not be reconstructed by dividing the two archive totals. That division gives a distance-weighted summary, while Model 1 gives every sequence equal weight. Both summaries are useful for different reporting questions, but they have different denominators. Keeping this distinction visible prevents an aggregation choice from becoming another reproducibility discrepancy layered onto the geometric mistake.

The coordinate-offset result is particularly informative because no physical movement was added. Correct camera-center distances remain unchanged, while the incorrect extraction changes by thousands of meters in aggregate. The finding therefore diagnoses a representation-dependent metric rather than scale recovery failure, extra camera motion, or sensitivity to a real translation of the camera through its environment.

7. A reproduction record needs more than array shape

The source rows all have the expected dimensions, and both inverse and forward transformations can be numerically valid. Those checks alone cannot detect that the wrong vector is being interpreted as a camera center. A reproduction record should additionally identify the source frame, destination frame, coordinate units, matrix multiplication convention, and geometric object extracted for each metric. A three-by-four shape is insufficient metadata for that task.

Influence analysis Cook, 1977 reminds us that a short sequence can have considerable weight in an equal-sequence average. We do not delete a sequence because its ratio is large or choose an origin offset to maximize the discrepancy. The one hundred meter coordinate change is a declared diagnostic, not an estimated physical parameter or a recommended stress-test distribution. Different offsets would create different naive errors while leaving the correctly measured paths invariant.

NumPy Harris et al., 2020 supports the array calculations, and the retained sequence table binds each reported ratio to its three path sums. Dataset documentation Gebru et al., 2021 and FAIR principles Wilkinson et al., 2016 motivate preserving this provenance alongside source checksums. Technical-debt concerns Sculley et al., 2015 further explain why an implicit transformation direction can become difficult to recover after results are copied into downstream reports.

The experiment does not estimate the prevalence of this failure among researchers or software packages. Its incorrect rule is intentionally constructed, and its large errors should not be presented as evidence that existing motion-transfer results are unreliable. Its value is narrower: an explicit coordinate-invariance check detects a concrete interpretation failure that file-integrity checks and a correct matrix inversion can both miss.

8. Why a simple translation test can miss the error

If orientation remains constant throughout a trajectory, the inverse translation is a common orthogonal transformation of the camera-center coordinates, with a sign change. Such a transformation preserves distances between consecutive positions. The naive extractor can therefore return the correct path length in that special case despite using the wrong geometric object. This is a mathematical counterexample to the idea that passing a pure-translation toy check certifies the extractor generally.

Changing orientation removes that protection. Each inverse translation then expresses the world origin in a different camera basis. Distances between those vectors can reflect the changing basis as well as the camera-center path. A common coordinate offset becomes orientation-dependent for exactly the same reason. These geometric observations explain the diagnostic design; they are not additional simulated trajectories included in the empirical sample.

The distinction is also relevant when a camera rotates around a fixed center. Correct translational travel can remain zero while an inverse-column vector changes if the chosen world origin is elsewhere. A reproduction check based on a single convenient origin might conceal this issue. We therefore combine the algebraic round trip with an origin transformation on the actual arrays, checking both that the correct extraction remains stable and that the deliberately wrong one need not do so.

These checks target meaning rather than just numerical execution. A program can be deterministic, match its saved checksum, and reproduce the wrong path sum exactly. Computational reproducibility establishes that the same operations were performed; geometric validation establishes whether those operations measure the intended quantity. The two are complementary requirements, and neither should be treated as a substitute for the other.

9. Conclusion

Inverse-pose translation and world camera position describe different vectors. Confusing them can greatly overstate path distance and introduce sensitivity to a physically irrelevant coordinate origin. The correct round trip preserves camera centers and path length to numerical precision. Reproducible camera-motion measurement therefore requires explicit transformation directions and geometric extraction rules, in addition to source files and computational code.

References

1. Triggs, Bill; McLauchlan, Philip F.; Hartley, Richard I.; Fitzgibbon, Andrew W.. (2000). Bundle Adjustment - A Modern Synthesis. Lecture Notes in Computer Science. https://doi.org/10.1007/3-540-44480-7_21.
2. Wilson, Benjamin; Qi, William; Agarwal, Tanmay; Lambert, John; Singh, Jagjeet; Khandelwal, Siddhesh; Pan, Bowen; Kumar, Ratnesh; Hartnett, Andrew; Pontes, Jhony Kaesemodel; Ramanan, Deva; Carr, Peter; Hays, James. (2023). Argoverse 2: Next Generation Datasets for Self-Driving Perception and Forecasting. arXiv:2301.00493. https://arxiv.org/abs/2301.00493.
3. Lourakis, Manolis I. A.; Argyros, Antonis A.. (2009). SBA: A software package for generic sparse bundle adjustment. ACM Transactions on Mathematical Software. https://doi.org/10.1145/1486525.1486527.
4. Schonberger, Johannes L.; Frahm, Jan-Michael. (2016). Structure-from-Motion Revisited. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). https://doi.org/10.1109/cvpr.2016.445.
5. Mur-Artal, Raul; Tardos, Juan D.. (2016). ORB-SLAM2: an Open-Source SLAM System for Monocular, Stereo and RGB-D Cameras. https://arxiv.org/abs/1610.06475.
6. Lepetit, Vincent; Moreno-Noguer, Francesc; Fua, Pascal. (2009). EPnP: An Accurate O(n) Solution to the PnP Problem. International Journal of Computer Vision. https://doi.org/10.1007/s11263-008-0152-6.
7. Yohann, Cabon,; Naila, Murray,; Martin, Humenberger,. (2020). Virtual KITTI 2. https://arxiv.org/abs/2001.10773.
8. Bland, J. Martin, and Douglas G. Altman. "Statistical Methods for Assessing Agreement between Two Methods of Clinical Measurement." *The Lancet*, vol. 327, no. 8476, 1986, pp. 307-310. https://doi.org/10.1016/S0140-6736(86)90837-8.
9. Hartley, Richard; Zisserman, Andrew. (2004). Multiple View Geometry in Computer Vision. Cambridge University Press. https://doi.org/10.1017/cbo9780511811685.
10. Gaidon, Adrien; Wang, Qiao; Cabon, Yohann; Vig, Eleonora. (2016). Virtual Worlds as Proxy for Multi-Object Tracking Analysis. CVPR 2016, 4340-4349. https://openaccess.thecvf.com/content_cvpr_2016/html/Gaidon_Virtual_Worlds_as_CVPR_2016_paper.html.
11. Higham, Nicholas J.. (2002). Accuracy and Stability of Numerical Algorithms: Second Edition. Society for Industrial and Applied Mathematics. https://doi.org/10.1137/1.9780898718027.
12. White, Halbert. “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity.” *Econometrica*, vol. 48, no. 4, 1980, pp. 817-838. https://doi.org/10.2307/1912934.
13. Geiger, A.; Lenz, P.; Urtasun, R.. (2012). Are we ready for autonomous driving? The KITTI vision benchmark suite. 2012 IEEE Conference on Computer Vision and Pattern Recognition. https://doi.org/10.1109/cvpr.2012.6248074.

14. Caesar, Holger, et al. "nuScenes: A multimodal dataset for autonomous driving" *arXiv preprint arXiv:1903.11027*, 2019. https://arxiv.org/abs/1903.11027.

15. Gebru, Timnit, et al. "Datasheets for datasets" *Communications of the ACM*, 2021. https://doi.org/10.1145/3458723.

16. Lin, Lawrence I-Kuei. (1989). A Concordance Correlation Coefficient to Evaluate Reproducibility. Biometrics. https://doi.org/10.2307/2532051.

17. Goldberg, David. (1991). What every computer scientist should know about floating-point arithmetic. ACM Computing Surveys. https://doi.org/10.1145/103162.103163.

18. Fisher, N. I.. (1993). Statistical Analysis of Circular Data. Cambridge University Press. https://doi.org/10.1017/cbo9780511564345.

19. Clement, Lee. (2026). pykitti: odometry.py. Source-code snapshot accessed7November2026; capture year, not original release year. https://github.com/utiasSTARS/pykitti/blob/master/pykitti/odometry.py.

20. Stevens, S. S.. (1946). On the Theory of Scales of Measurement. Science. https://doi.org/10.1126/science.103.2684.677.

21. Shoemake, Ken. (1985). Animating rotation with quaternion curves. Proceedings of the 12th annual conference on Computer graphics and interactive techniques - SIGGRAPH '85. https://doi.org/10.1145/325334.325242.

22. Cameron, A. Colin, and Douglas L. Miller. “A Practitioner’s Guide to Cluster-Robust Inference.” *Journal of Human Resources*, vol. 50, no. 2, 2015, pp. 317-372. https://doi.org/10.3368/jhr.50.2.317.

23. Cook, R. Dennis. (1977). Detection of Influential Observation in Linear Regression. Technometrics. https://doi.org/10.1080/00401706.1977.10489493.

24. Harris, Charles R.; Millman, K. Jarrod; van der Walt, Stéfan J.; Gommers, Ralf; Virtanen, Pauli; Cournapeau, David; Wieser, Eric; Taylor, Julian; Berg, Sebastian; Smith, Nathaniel J.; Kern, Robert; Picus, Matti; Hoyer, Stephan; van Kerkwijk, Marten H.; Brett, Matthew; Haldane, Allan; del Río, Jaime Fernández; Wiebe, Mark; Peterson, Pearu; Gérard-Marchant, Pierre; Sheppard, Kevin; Reddy, Tyler; Weckesser, Warren; Abbasi, Hameer; Gohlke, Christoph; Oliphant, Travis E.. (2020). Array programming with NumPy. Nature. https://doi.org/10.1038/s41586-020-2649-2.

25. Sun, Pei, et al. "Scalability in Perception for Autonomous Driving: Waymo Open Dataset" *arXiv preprint arXiv:1912.04838*, 2019. https://arxiv.org/abs/1912.04838.

26. Meng, Zijie. (2026). ParaScale: Scale-Calibrated Camera-Motion Transfer via a Gauge-Invariant Parallax Number. https://arxiv.org/abs/2606.19805.

27. Wilkinson, Mark D.; Dumontier, Michel; Aalbersberg, IJsbrand Jan; Appleton, Gabrielle; Axton, Myles; Baak, Arie; Blomberg, Niklas; Boiten, Jan-Willem; da Silva Santos, Luiz Bonino; Bourne, Philip E.; Bouwman, Jildau; Brookes, Anthony J.; Clark, Tim; Crosas, Mercè; Dillo, Ingrid; Dumon, Olivier; Edmunds, Scott; Evelo, Chris T.; Finkers, Richard; Gonzalez-Beltran, Alejandra; Gray, Alasdair J.G.; Groth, Paul; Goble, Carole; Grethe, Jeffrey S.; Heringa, Jaap; ’t Hoen, Peter A.C; Hooft, Rob; Kuhn, Tobias; Kok, Ruben; Kok, Joost; Lusher, Scott J.; Martone, Maryann E.; Mons, Albert; Packer, Abel L.; Persson, Bengt; Rocca-Serra, Philippe; Roos, Marco; van Schaik, Rene; Sansone, Susanna-Assunta; Schultes, Erik; Sengstag, Thierry; Slater, Ted; Strawn, George; Swertz, Morris A.; Thompson, Mark; van der Lei, Johan; van Mulligen, Erik; Velterop, Jan; Waagmeester, Andra; Wittenburg, Peter; Wolstencroft, Katherine; Zhao, Jun; Mons, Barend. (2016). The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data. https://doi.org/10.1038/sdata.2016.18.

28. Sculley, D., et al. "Hidden Technical Debt in Machine Learning Systems." *Advances in Neural Information Processing Systems*, vol. 28, 2015. https://papers.nips.cc/paper/2015/hash/86df7dcfd896fcaf2674f757a2463eba-Abstract.html

29. Carroll, Raymond J., David Ruppert, Leonard A. Stefanski, and Ciprian M. Crainiceanu. “Measurement Error in Nonlinear Models.” 2006. https://doi.org/10.1201/9781420010138.

30. Virtanen, Pauli; Gommers, Ralf; Oliphant, Travis E.; Haberland, Matt; Reddy, Tyler; Cournapeau, David; Burovski, Evgeni; Peterson, Pearu; Weckesser, Warren; Bright, Jonathan; van der Walt, Stéfan J.; Brett, Matthew; Wilson, Joshua; Millman, K. Jarrod; Mayorov, Nikolay; Nelson, Andrew R. J.; Jones, Eric; Kern, Robert; Larson, Eric; Carey, C J; Polat, İlhan; Feng, Yu; Moore, Eric W.; VanderPlas, Jake; Laxalde, Denis; Perktold, Josef; Cimrman, Robert; Henriksen, Ian; Quintero, E. A.; Harris, Charles R.; Archibald, Anne M.; Ribeiro, Antônio H.; Pedregosa, Fabian; van Mulbregt, Paul; SciPy 1.0 Contributors; Vijaykumar, Aditya; Bardelli, Alessandro Pietro; Rothberg, Alex; Hilboll, Andreas; Kloeckner, Andreas; Scopatz, Anthony; Lee, Antony; Rokem, Ariel; Woods, C. Nathan; Fulton, Chad; Masson, Charles; Häggström, Christian; Fitzgerald, Clark; Nicholson, David A.; Hagen, David R.; Pasechnik, Dmitrii V.; Olivetti, Emanuele; Martin, Eric; Wieser, Eric; Silva, Fabrice; Lenders, Felix; Wilhelm, Florian; Young, G.; Price, Gavin A.; Ingold, Gert-Ludwig; Allen, Gregory E.; Lee, Gregory R.; Audren, Hervé; Probst, Irvin; Dietrich, Jörg P.; Silterra, Jacob; Webber, James T; Slavič, Janko; Nothman, Joel; Buchner, Johannes; Kulick, Johannes; Schönberger, Johannes L.; de Miranda Cardoso, José Vinícius; Reimer, Joscha; Harrington, Joseph; Rodríguez, Juan Luis Cano; Nunez-Iglesias, Juan; Kuczynski, Justin; Tritz, Kevin; Thoma, Martin; Newville, Matthew; Kümmerer, Matthias; Bolingbroke, Maximilian; Tartre, Michael; Pak, Mikhail; Smith, Nathaniel J.; Nowaczyk, Nikolai; Shebanov, Nikolay; Pavlyk, Oleksandr; Brodtkorb, Per A.; Lee, Perry; McGibbon, Robert T.; Feldbauer, Roman; Lewis, Sam; Tygier, Sam; Sievert, Scott; Vigna, Sebastiano; Peterson, Stefan; More, Surhud; Pudlik, Tadeusz; Oshima, Takuya; Pingel, Thomas J.; Robitaille, Thomas P.; Spura, Thomas; Jones, Thouis R.; Cera, Tim; Leslie, Tim; Zito, Tiziano; Krauss, Tom; Upadhyay, Utkarsh; Halchenko, Yaroslav O.; Vázquez-Baeza, Yoshiki. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods. https://doi.org/10.1038/s41592-019-0686-2.